PARALLEL COMPUTATION OF SENSITIVITY DERIVATIVES WITH APPLICATION TO AERODYNAMIC OPTIMIZATION OF A WING

Robert T. Biedron (*r.t.biedron@larc.nasa.gov 757-864-2156*)
Jamshid A. Samareh (*j.a.samareh@larc.nasa.gov* 757-864-5776)
Lawrence L. Green (*l.l.green@larc.nasa.gov 757-864-2228*)

NASA Langley Research Center
Hampton, VA 23681-2199

ABSTRACT

This paper focuses on the parallel computation of aerodynamic derivatives via automatic differentiation of the Euler/Navier-Stokes solver CFL3D. The comparison with derivatives obtained by finite differences is presented and the scaling of the time required to obtain the derivatives relative to the number of processors employed for the computation is shown. Finally, the derivative computations are coupled with an optimizer and surface/volume grid deformation tools to perform an optimization to reduce the drag of a three-dimensional wing.

INTRODUCTION

Recently researchers have shown a great deal of interest in the application of advanced CFD methods to aerodynamic optimization, for both single-discipline and multidiscipline applications. Central to any gradient-based aerodynamic optimization problem is the evaluation of solution derivatives with respect to the chosen design variables. Differentiation of the CFD source code used to obtain the solution gives *exact* derivatives of the discrete equations, without the step size problems of finite differences. Although quite tedious to perform by hand, exact differentiation of a source code is readily accomplished using an automatic differentiation (AD) tool such as ADIFOR[1]. The computational time for AD derivatives scales with the number of design variables, and the computational time may be prohibitive for large number of design variables. One way to reduce the effective computation time (wall time) is to subdivide the computational domain and compute each subdomain on a different processor. For this approach to be useful, the computational code must scale well with increasing number of processors.

The CFD code used for this study, CFL3D[2], has been widely used for aerodynamic analysis on complex configurations. One version of the code (CFL3Dv4.1hp) has recently been ported to parallel computer architecture via the use of MPI protocols. Studies have indicated good scaling on Origin 2000 testbeds for Euler and Navier-Stokes solutions[3]. Even more recently the parallel code has been passed through the ADIFOR automatic differentiation tool to generate code capable of computing both the solution and the gradient of the solution with respect to geometric design variables.

PARAMETERIZATION AND DESIGN VARIABLES

For aerodynamic optimization, a parameterized surface definition that relates the shape to geometric design variables is required. In many instances, a computational grid defining the baseline shape of the configuration is readily available, but a parameterization of the surface is not.

A surface parameterization scheme that overcomes this difficulty has recently been developed by the second author[4]. The method is a free-form deformation approach very similar to morphing techniques used in computer animation. It can simulate planform, twist, dihedral, thickness, and camber variations. In a sense, the model is treated as putty or clay in areas where it can be twisted, bent, tapered, compressed or expanded, but retains the same topology. The method is equally applicable to computational structures grids, and thus is ideally suited for aerostructural calculations.

An existing grid defining the ONERA M6 wing[5] was parameterized with 52 parameters. Of those 52 parameters, 31 were chosen as design variables: 5 planform, 4 twist, 4 shear, 9 thickness and 9 camber. Figure 1 shows the locations of the design variables chosen for the wing optimization.

COMPARISON WITH FINITE DIFFERENCES

As a validation that the AD code produces the correct derivatives, comparisons with central finite differences (FD) were made using double-precision arithmetic. The AD derivatives and finite differences were computed for inviscid flow over the M6 wing with the design variables described above. The flow conditions were Mach 0.84 and $\alpha = 3.06°$. A coarse grid of dimensions 97x17x17 in the streamwise, spanwise and normal directions, respectively, was used for the derivative validation studies. For the FD results, residuals were driven to machine zero; for the AD results, computations were stopped when derivatives no longer varied in the fifth decimal place. All finite differences were computed using a step size of 10^{-6}. Experience has shown that single precision is sufficient for inviscid analysis, e.g. negligible difference in force coefficients between single and double precision. The AD calculations were repeated with single precision to see if the same would hold true for derivatives of the force coefficients.

The results are summarized in Table 1 for several representative design variables. Similar results are obtained for other derivatives and other force coefficients. It is evident that the AD code does in fact produce the correct derivatives. Note that the AD results are the same for both single and double precision, at least to 4 decimal places, a result typical of other derivatives as well. Thus, the AD code can be used reliably with single precision, at least for the inviscid flow considered here. The advantage is that the code runs approximately 40% faster in single precision. Although not shown, it should be noted that when used with single precision, finite differences could not be obtained with better than approximately one percent error as compared to double precision. Furthermore, different design variables required different step sizes to obtain even that level of accuracy.

SCALING STUDY

The scaling study was carried out for inviscid flow over a High Speed Civil Transport (HSCT) configuration at Mach 2.4. The grid used was comprised of approximately 540,000 cells in 64 equal sized blocks. The surface was parameterized with 27 design variables in a manner similar to that used for the M6 wing described above. For the scaling study, 100 three-level multigrid iterations were used, resulting in derivatives that remained unchanging with iteration number through the fifth decimal place. The computations were carried out in single precision. The results were obtained on Origin 2000 computers, using from 1 to 32 compute processors (an extra processor functions as the host, performing I/O tasks which consume relatively little CPU time). Each case was run at least twice to try to account for run-to-run variations due to system load.

The scaling results are shown in Figure 2. Computing only the solution (no derivatives) required 0.787 hours on a single processor, dropping to 0.023 hours on 32 processors. Computing the 27

AD derivatives along with the solution required nearly 33.5 hours on a single processor, dropping to 1.05 hours on 32 processors. The speedup was essentially linear for both solution and solution plus gradient calculations.

WING OPTIMIZATION

As an application of the parallel AD code, an aerodynamic optimization of an ONERA M6 wing was carried out. The objective of the optimization was to minimize the drag while maintaining the same lift as the baseline design. As for the derivative validation, inviscid flow at Mach 0.84 and $\alpha = 3.06°$ was used, however a finer grid of dimensions 197x33x33 was employed for the optimization. The design variables used were the 31 shown in Figure 1, although for the current study the planform variables were constrained so that they did not change during the design, resulting in a fixed wing area. This eliminated the need for an additional code to calculate the wing area and derivatives of the wing area with respect to the design variables. Also, to prevent negative cell volumes near the tip, thickness variables Th3, Th6, and Th9 were constrained so as not to change. Design variable limits were arbitrarily chosen as follows: twist, +/- $1°$; all others, +/- 1 percent span.

The optimizer used for this work is a modified version of the CONMIN code[6] known as JOPT[7]. Within each optimization cycle, the solution and gradient data provided to the optimizer are used to determine a linear approximation to the objective function and constraints used in the 1D line searches. This makes each line search much faster, but the linear approximation is only valid with a small region of the current solution. User-defined move limits for the design variables are required to insure that the optimizer searched only where the current linear approximation was reasonable.

The solution and design-variable changes suggested by the optimizer were incorporated into the surface model using the geometry deformation scheme mentioned earlier. Next, an AD version of the CSCMDO code[8] was used to propagate the difference between the old and new surfaces smoothly throughout the volume grid, determining the grid sensitivities in the process.

Figure 3 shows the design cycle history for both lift and drag. In this optimization, the angle of attack is fixed, and it was found that in order to move away from the current design, the constraint on the lift coefficient had to be relaxed temporarily. This is shown clearly in the figure: for the first 19 design cycles, C_L is allowed to deviate by up to 0.01 from the desired value. After design cycle 19 the tolerance on the lift constraint is tightened to 10^{-6}. The drag increased slightly when the lift constraint was tightened, but after the initial rise there was no further change in drag at the target lift coefficient. The net result was approximately 29 counts of drag reduction at the baseline lift. Figures 4 and 5 show comparisons of the solutions computed on the initial and final designs. The results indicate a significant reduction in the shock strength at most spanwise stations. Also shown in Figure 5 are initial and final wing sections at selected spanwise stations.

Using 16 compute processors on a 250 Mhz Origin 2000, each design cycle took approximately 115 minutes, of which approximately 100 minutes was spent in evaluating the 31 gradients, using 300 multigrid cycles. The time per design cycle can be reduced as desired by increasing the number of processors employed. Although not done in this preliminary study, it should be possible to further reduce the total optimization cost by utilizing the mesh sequencing option in CFL3D to perform most of the design variable changes on a coarser level, and only then moving up to the finest level for the final design cycles.

CONCLUDING REMARKS

A parallel, differentiated version of the CFL3D code has been demonstrated to yield accurate derivatives with respect to geometric design variables. Furthermore, these computationally intensive derivative calculations have been shown to scale well with increasing number of processors. The parallel AD code was coupled to grid deformation and optimization packages and used to reduce the inviscid, transonic drag on a wing. Future applications will consider viscous flows.

REFERENCES

1. Bischof, C., Carle, A., Khademi, P., Mauer, A., and Hovland, P.; "ADFIFOR 2.0 User's Guide," Argonne National Laboratory Mathematics and Computer Science Division Technical Memorandum No. 192, August, 1995.

2. Krist, S. L., Biedron, R. T., and Rumsey, C. L.; "CFL3D Users Manual (Version 5.0)," NASA TM-1998-208444, June 1998.

3. Faulkner, T.; "Origin 2000 Update: Studies Show CFL3D Can Obtain Reasonable Performance," NAS News, Vol. 2, No. 27, Nov.-Dec. 1997.

4. Samareh, J. A.; "Geometry Modeling and Grid Generation for Design and Optimization," ICASE/LaRC/NSF/ARO Workshop On Computational Aerosciences In The 21st Century, Hampton, VA, April 22-24, 1998.

5. Schmitt, V., and Charpin, F.; "Pressure Distributions on the ONERA-M6 Wing at Transonic Mach Numbers," Experimental Data Base for Computer Program Assessment, AGARD-AR-138, May 1979, pp. B1-1 - B1-41.

6. Vanderplaats, G. N.; "CONMIN- A FORTRAN Program For Constrained Function Minimization," NASA TM X-62282, August, 1973.

7. Walsh, J. L., Young, K. C., Pritchard, J. I., Adelman, H. M., and Mantay, W. R.; "Integrated Aerodynamic/Dynamic/Structural Optimization of Helicopter Rotor Blades Using Multilevel Decomposition,", NASA TP 3465 / ARL TR 518, January, 1995.

8. Jones, W. T., and Samareh-Abolhassani, J.; "A Grid Generation System for Multi-disciplinary Design Optimization," Proceedings from the 12th AIAA Computational Fluid Dynamics Conference, AIAA-95-1689, San Diego, CA, June 1995, pp. 657-669.

Derivative	AD (DP)	FD (DP)	% error (DP)	AD (SP)
$dC_L/d(Tw\ 3)$	-0.02944	-0.02944	0.0	-0.02944
$dC_L/d(Th\ 8)$	+0.43321	+0.43321	0.0	+0.43323
$dC_L/d(Ca\ 8)$	+2.8380	+2.8380	0.0	+2.8380
$dC_D/d(Tw3)$	-0.00246	-0.00246	0.0	-0.00246
$dC_D/d(Th\ 8)$	+0.07016	+0.07016	0.0	+0.07016
$dC_D/d(Ca\ 8)$	+0.16467	+0.16467	0.0	+0.16467

Table 1. Accuracy of lift and drag coefficient derivatives computed using automatic differentiation and central finite differences. DP denotes double precision; SP denotes single precision.

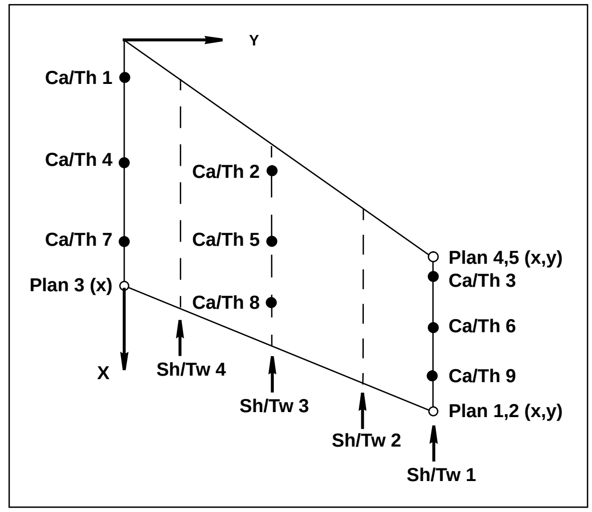

Figure 1. Design variable locations; Ca/Th denotes camber/thickness variables at points indicated by the solid circles; Sh/Tw denotes shear/twist variables, defined along the dashed lines; Plan denotes planform variables, at points indicated by the empty circles.

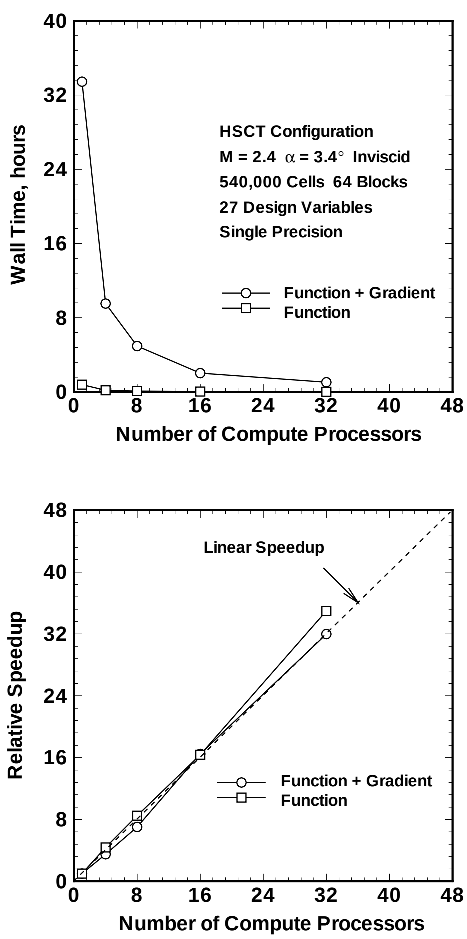

Figure 2. Origin 2000 scaling for both solution and solution plus gradient evaluation for an HSCT configuration with 27 design variables.

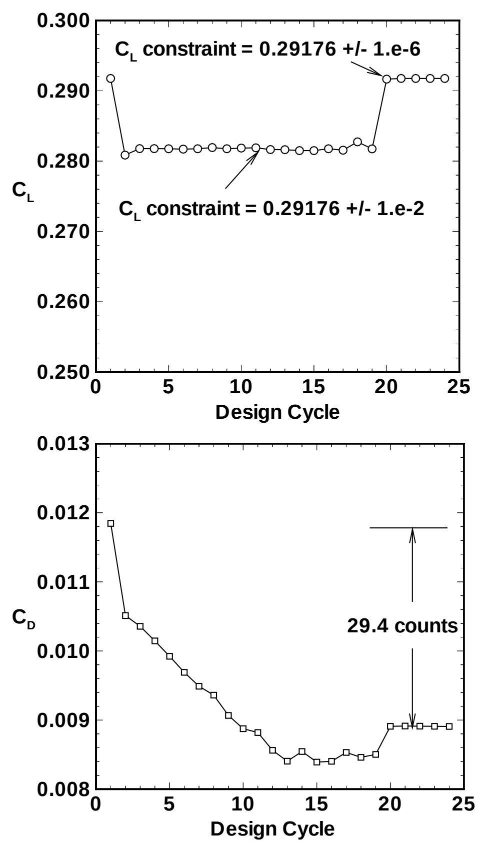

Figure 3. Design cycle history for ONERA M6 wing optimization.

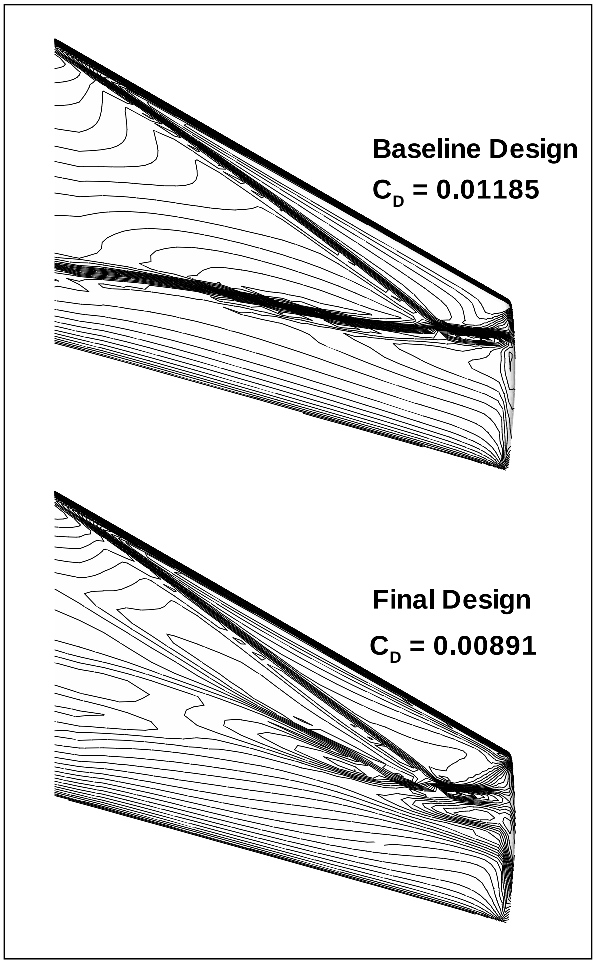

Figure 4. Comparison of surface pressures on the final wing design with the baseline M6 wing.

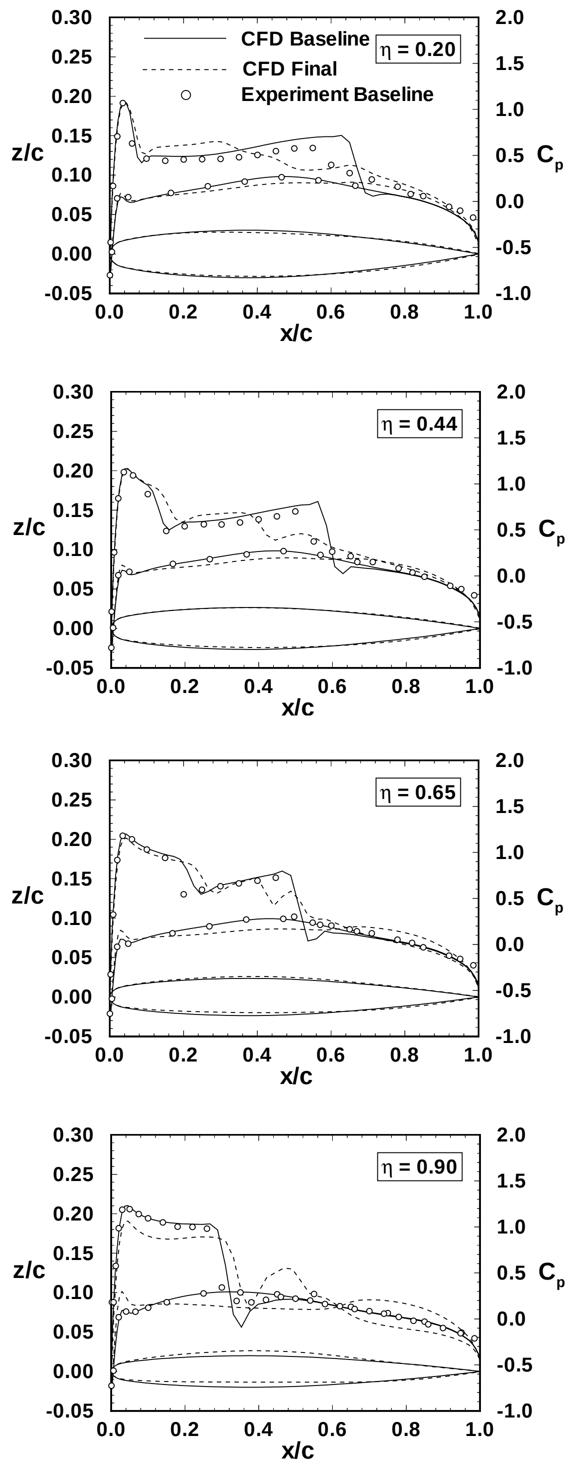

Figure 5. Comparison of initial and final Cp distribution and wing cross section at selected spanwise stations.